100 NOC Engineer Interview Scenarios

Scenario-based interview questions and model answers for Network Operations Center engineers.

Routing & Switching

Q1. A user can't reach a server in another VLAN, but can reach one in the same VLAN. What do you check first?

A: Verify a routed path exists between VLANs (SVI/L3 route), check ACLs on the L3 switch/router, and confirm the default gateway on the client's subnet is correct.

Q2. Two switches show a spanning-tree topology change every few minutes. What's happening?

A: Likely a flapping link or duplex mismatch causing repeated TCNs. Check port error counters and interface logs for up/down flaps, and verify STP port roles aren't oscillating.

Q3. After adding a new switch to the network, users report intermittent broadcast storms. Cause?

A: A loop was created — likely the new switch's uplink wasn't added to the STP topology correctly, or a port was misconfigured (trunk vs access), creating a physical loop without loop protection (BPDU guard / loop guard disabled).

Q4. A route to a remote site works via traceroute but pings drop 50% of the time. What's suspect?

A: Possible asymmetric routing or a flaky link with intermittent packet loss. Check interface error counters and compare inbound/outbound paths for asymmetry.

Q5. OSPF neighbors won't form between two routers on the same subnet. First checks?

A: Confirm matching area IDs, hello/dead timers, MTU, and authentication settings. Verify they're on the same subnet and no ACL is blocking OSPF multicast (224.0.0.5/6).

Q6. A newly added VLAN isn't passing traffic across a trunk link. Likely cause?

A: VLAN isn't allowed on the trunk (needs adding to the allowed VLAN list) or isn't created/active on one of the switches. Also check for a native VLAN mismatch.

Q7. Users on one floor lose connectivity randomly, correlating with cleaning staff hours. Suspected cause?

A: A physical cable being disturbed or unplugged. Check interface flap timestamps against the reported times and inspect/reseat the cabling.

Q8. Router CPU spikes to 99% during peak hours, causing slow routing. What do you check?

A: Identify the process consuming CPU ('show processes cpu'), check for high traffic volume, heavy ACL processing, routing table churn, or a DoS/scan. Also confirm hardware/CEF switching hasn't fallen back to software switching.

Q9. A static route works for one subnet but not its neighbor subnet on the same interface. Cause?

A: A missing or incorrect static route/subnet mask for the second subnet. Verify routing table entries for the exact prefix and mask.

Q10. Interface counters show high input errors and CRC errors. Root cause?

A: A physical layer issue — duplex mismatch, a bad cable, or a bad SFP/interference. Check that duplex/speed settings match on both ends and replace the suspect cable or transceiver.

Q11. Multiple users behind an EtherChannel report slow throughput despite it showing "up." What to check?

A: Verify all member links are actually up and that the load-balancing hash isn't sending all traffic down a single link. Check for one leg with errors reducing effective bandwidth.

Q12. A default route was removed accidentally and internet access failed network-wide. How do you prevent recurrence?

A: Require change control/peer review for routing changes, use route redundancy (floating static or dynamic routing), and keep configuration backups/versioning for fast rollback.

DNS & DHCP

Q13. Users can reach IPs directly but not by hostname. What's the fix path?

A: Confirm DNS server reachability (ping/nslookup), check the client's DNS settings, then test resolution directly against the DNS server to isolate a client vs. server issue.

Q14. Internal DNS resolves fine, but external domains fail intermittently. Cause?

A: A forwarder/upstream DNS issue, or a firewall blocking port 53 (especially TCP for larger responses). Check forwarder health and DNS server logs.

Q15. A new laptop can't get an IP address on the network. Where do you look?

A: Check the DHCP scope for available leases, verify the DHCP relay/helper address on the VLAN's gateway, and confirm the port isn't blocked by port security or 802.1x.

Q16. A branch office DHCP scope is exhausted despite low headcount. Cause?

A: Stale leases from disconnected devices not being reclaimed, or a lease time that's too short. Check the lease table and adjust lease duration or clear stale entries.

Q17. A DNS server is up, but responses for one specific domain always time out. Cause?

A: That domain's authoritative servers may be unreachable, or a zone forwarder is misconfigured. Test with dig/nslookup directly against the authoritative server.

Q18. Devices are getting APIPA (169.254.x.x) addresses. What's the first thing to check?

A: The DHCP server or relay is unreachable — check DHCP service status, the relay/helper configuration on the router, and VLAN trunking to the DHCP server.

Q19. A dual-homed server intermittently registers the wrong IP in DNS. Cause?

A: Dynamic DNS registration racing between two NICs. Disable registration on the non-primary interface or set the correct interface priority/metric.

Q20. After a DNS record change, some users still resolve the old IP. Why?

A: DNS caching — the client resolver cache or a downstream/ISP DNS cache hasn't expired yet due to TTL. Flush the local cache and wait for TTL expiration.

Firewall & Security

Q21. Traffic is allowed by an ACL but still blocked. What's next?

A: Check for an earlier deny rule matching first (ACL order matters), verify NAT/routing, and check for additional security policies such as zone-based firewall or stateful inspection.

Q22. After a firewall rule update, VoIP calls drop after 30 seconds. Cause?

A: Likely a UDP session timeout or a missing ALG/SIP inspection setting. Adjust the UDP/SIP session timeout or enable proper protocol inspection.

Q23. Users report they can't access a site that used to work, now blocked by URL filtering. Checklist?

A: Confirm the category caused the block (check filter logs), check for a recent category database update that reclassified the site, and rule out an SSL-inspection certificate issue.

Q24. A perimeter firewall shows a spike in denied connections from a single external IP. Response?

A: Investigate as a potential scan or attack — check the pattern (port scan, brute force), apply a temporary block/rate-limit, and escalate per the security incident process.

Q25. A site-to-site tunnel is up but no traffic passes. Diagnostic order?

A: Check Phase 1/Phase 2 SA status, verify interesting traffic/ACLs match the encryption domain, confirm routing sends traffic into the tunnel, and check for NAT-T issues.

Q26. An internal server suddenly can't be reached from outside after a change. First check?

A: Confirm the NAT/port-forward rule is intact, the firewall policy allows the traffic, and the server itself is actually listening (not a service-down issue).

Q27. A user's laptop is quarantined by NAC without explanation. Steps to diagnose?

A: Check NAC policy/compliance logs for a failed posture check (AV, patch level), verify certificate/802.1x supplicant configuration, and review the specific failure reason in the NAC console.

Q28. Logs show repeated login failures from many source IPs against one admin account. Response?

A: Treat it as brute-force/credential stuffing — lock or reset the account, verify account lockout and MFA are enabled, block the offending IPs, and escalate to security.

Q29. After enabling deep packet inspection, throughput drops. What's the trade-off and diagnostic?

A: DPI is CPU-intensive on the firewall. Check firewall CPU/session load, consider more selective inspection policies or hardware acceleration, and benchmark before/after.

Q30. A misconfigured ACL blocked a critical application at 2 AM. What process failure led to this, and how do you prevent it?

A: A lack of change control/testing window and rollback plan. Require peer review, staged deployment, and a pre-approved rollback for any firewall/ACL change.

Q31. Users can browse sites but file downloads consistently fail or hang. Cause?

A: A possible MTU/fragmentation issue interacting with the firewall, or IPS blocking large-payload patterns. Check TCP MSS clamping and IPS logs.

Q32. Remote sites behind different firewall vendors report inconsistent VPN client compatibility. Cause?

A: A firmware version mismatch, or differing NAT-T/UDP 500-4500 handling. Standardize firmware and verify NAT traversal support across all peer devices.

VPN & Remote Access

Q33. A remote worker connects to VPN but can't access internal file shares. What to check?

A: Verify VPN split-tunnel routing includes the internal subnets, DNS resolves to the internal DNS server while connected, and firewall rules permit SMB traffic from the VPN pool.

Q34. A VPN client repeatedly disconnects every ~1 hour. Cause?

A: Likely an IKE/SA lifetime mismatch or an idle timeout. Align rekey timers on both ends and check NAT keepalive settings.

Q85. You suspect another team's recent change caused an issue, but they deny it. How do you proceed?

A: Present objective evidence — logs and timestamps correlating the change to the symptom onset — and request a joint troubleshooting session rather than assigning blame unilaterally.

Q86. How do you triage several simultaneous P2 incidents with limited staff on shift?

A: Assess business impact and affected user count for each, tackle the highest-impact one first, and communicate realistic timelines for the others while pulling in extra resources if available.

Q87. A recurring issue keeps getting temporary fixes but never a permanent resolution. How do you address this?

A: Push for formal problem management (ITIL) — open a problem record to investigate the root cause separately from the incident-by-incident firefighting.

Q88. How would you explain a complex network outage to a non-technical stakeholder?

A: Use plain-language analogies (e.g., "a blocked road" instead of "routing loop"), focus on impact and ETA rather than technical detail, and avoid jargon.

Q89. A change you made during a maintenance window caused unexpected downtime. What's your immediate next step?

A: Roll back the change immediately per the pre-approved rollback plan, confirm service restoration, then document and report the incident transparently.

Q90. How do you keep yourself and the team updated on network changes happening across departments?

A: Actively review the change calendar/CAB notifications, maintain a shift handoff log, and encourage a culture of proactive change communication.

Server & OS Networking

Q91. A Linux server can ping its gateway but not any other host. Cause?

A: Possible missing or incorrect routing table entries beyond the local subnet, or a host-based firewall (iptables/firewalld) blocking forwarded traffic.

Q92. A Windows server loses network connectivity right after a Windows Update. Cause?

A: A NIC driver issue introduced by the update, or the network profile reset to "Public," blocking previously allowed traffic. Check the driver version and firewall profile.

Q93. Two NICs on a server on the same subnet cause intermittent connectivity. Cause?

A: Routing table ambiguity or a binding-order issue with multiple NICs on the same subnet. Disable one NIC or reconfigure to avoid dual-homing on the same subnet.

Q94. A server's network throughput is capped well below the NIC's rated speed. Cause?

A: A duplex/speed mismatch, driver/offload settings, or a vSwitch bottleneck in a virtualized environment. Check the negotiated speed and vSwitch configuration.

Q95. A server's hostname resolves fine internally, but SSL certificate errors happen intermittently. Cause?

A: A mismatch between DNS load balancing across multiple servers with different certificates, or a stale cert. Verify the certificate SANs match all resolvable hostnames.

Misc Protocols & General

Q96. Users complain email is delayed by up to an hour but eventually delivers. Cause?

Q73. One-way audio is reported on VoIP calls between two sites. Cause?

A: A classic NAT/firewall/SIP ALG issue affecting the RTP stream in one direction. Check SIP ALG settings and the RTP port range through the firewalls.

Q74. Softphone calls drop exactly every 3 minutes. Cause?

A: A possible SIP registration/re-INVITE timeout, or a NAT session timeout shorter than the call's keep-alive interval. Adjust the NAT timeout or registration refresh interval.

Q75. VoIP quality complaints only come from wireless users, not wired ones. Cause?

A: A lack of QoS/WMM on the WLAN, or excessive Wi-Fi retransmissions. Verify QoS marking is honored across the wireless controller.

Q76. After a WAN link upgrade to higher bandwidth, voice quality got worse, not better. Cause?

A: The QoS policy is likely still capped/shaped for the old bandwidth, or misclassifying traffic. Re-tune the QoS policy for the new link speed and confirm correct traffic classification.

Q77. Conference bridge calls degrade only when more than 10 participants join. Cause?

A: Media server/MCU resource exhaustion (CPU or bandwidth) at higher participant counts. Check media server capacity and consider scaling or codec optimization.

Physical Layer & Cabling

Q78. A newly terminated cable passes a link light but has a high error rate. Cause?

A: Poor termination — wire pairs not matching the TIA/EIA standard or untwisted too far — causing crosstalk. Re-terminate and test with a cable certifier.

Q79. A fiber link shows light-loss readings well beyond spec. Cause?

A: Dirty or damaged fiber connectors, excessive bends, or too many splices. Clean the connectors and test with an OTDR/light meter.

Q80. A patch panel port works, but the same run tested at the wall jack doesn't. Cause?

A: Possible physical damage to the cable run inside the wall, or a bad jack termination. Test each segment individually with a cable tester to isolate the fault.

Q81. A copper run works fine at 100Mbps but fails at gigabit speeds over the same cable. Cause?

A: A cable length/quality issue (e.g., Cat5 instead of Cat5e/6) or a wiring standard mismatch affecting the extra pairs gigabit needs. Verify cable category and termination.

Q82. An SFP transceiver shows link down despite being properly seated. Cause?

A: Mismatched SFP type/vendor lock-out, a wavelength mismatch (single-mode vs. multi-mode), or a faulty SFP. Verify compatibility and swap to test.

Q83. Multiple ports on the same switch card show intermittent errors after a power event. Cause?

A: Possible hardware damage from a surge/brownout. Check switch logs for hardware faults and consider replacing the affected line card or switch.

Escalation, ITIL & Soft Skills

Q84. A vendor's ticket for a circuit outage isn't progressing after 2 hours. What do you do?

A: Escalate within the vendor's own support tier using their escalation path, and open an internal bridge to keep stakeholders informed while pursuing the vendor's SLA commitments.

Q61. How would you handle an executive demanding constant updates during a major outage?

A: Assign a dedicated communicator separate from the engineers actively fixing the issue, and send updates at agreed intervals through the incident communication channel.

Q62. A change made by another team caused an outage but wasn't communicated to the NOC. Process improvement?

A: Enforce mandatory Change Advisory Board (CAB) notifications to the NOC before any network-impacting change, backed by a change calendar visible to the whole team.

Q63. Your shift is ending but an incident isn't resolved. How do you hand it off?

A: Give a thorough handoff: current status, what's already been tried, planned next steps, and any relevant vendor case numbers — all documented directly in the ticket.

Cloud & Load Balancing

Q64. A load balancer keeps sending traffic to a backend server that's actually down. Cause?

A: The health check is likely misconfigured — checking the wrong port/path, or polling too infrequently. Verify the health check settings and thresholds.

Q65. Users experience session drops with a round-robin load balancer on a stateful app. Cause?

A: A lack of session persistence. Enable sticky sessions (cookie-based or source-IP based) for stateful applications.

Q66. Cloud VPC instances can't reach the internet despite having a public IP. Cause?

A: A missing or misconfigured internet gateway/route table entry, or a security group/NACL blocking outbound traffic.

Q67. A hybrid cloud connection (Direct Connect/ExpressRoute) shows intermittent BGP flaps. Cause?

A: A physical layer issue on the circuit, or a BGP timer mismatch. Check circuit provider status and BGP session logs for hold-timer expiry.

Q68. An app in the cloud is reachable internally but not from on-prem over VPN. Cause?

A: A route propagation issue between the VPN gateway and VPC route tables, or overlapping CIDR ranges between on-prem and cloud. Verify route tables and check for IP conflicts.

Q69. Auto-scaling added new instances, but the load balancer isn't sending them traffic. Cause?

A: New instances aren't registered in the target group, or are failing health checks during startup. Check target group registration and the health check grace period.

Q70. Cross-region cloud traffic has unexpectedly high latency. Cause?

A: Traffic may be routing over the public internet instead of a private backbone/peering. Verify use of the cloud provider's private network/interconnect.

Q71. A cloud security group update blocked traffic that used to work through the NAT gateway. Cause?

A: The NAT gateway's outbound IP wasn't added to the security group's allow list after the change. Add or update the rule for the NAT IP range.

VoIP & QoS

Q72. Phone calls have good quality most of the day but become choppy around 5 PM. Cause?

A: Likely bandwidth contention from end-of-day traffic (backups, updates). Verify QoS is correctly prioritizing voice (DSCP EF marking) end-to-end.

Q48. Wi-Fi calling voice quality is choppy while data works fine. Cause?

A: A lack of QoS/WMM prioritization for voice traffic on the WLAN, or roaming delay between APs. Enable WMM and check for fast roaming (802.11r) support.

Monitoring, Alerts & NOC Process

Q49. Monitoring floods with hundreds of alerts during a single outage. How do you prioritize?

A: Identify the root-cause alert — usually the highest-level device that went down — and suppress/correlate the downstream dependent alerts. Follow the incident's impact/priority matrix.

Q50. A critical alert didn't page anyone during last night's outage. What do you investigate?

A: Check the alert routing/escalation policy configuration, the on-call schedule accuracy, and whether the alert threshold/severity mapping was correct.

Q51. You get a flapping "interface down/up" alert every 5 minutes on a core switch. What's your response?

A: Investigate the physical layer (cable/SFP) or a misbehaving attached device. If it's confirmed non-critical noise, adjust alert dampening/hold-down timers.

Q52. Ticket volume for the same issue is coming in from multiple sites. How do you handle it?

A: Treat it as a major/parent incident, link related tickets, send a single status update, and investigate a common root cause (shared upstream provider, central service, etc.).

Q53. Monitoring shows 100% packet loss to a device that users say is working fine. Cause?

A: The monitoring server's path to the device may be broken (ICMP blocked, routing issue) while the user traffic path is different. Verify from multiple vantage points.

Q54. How do you decide whether to escalate to Tier 2/3 vs. keep troubleshooting?

A: Follow the escalation matrix based on time-to-resolve SLAs and business impact severity, and escalate proactively once standard runbook steps are exhausted — don't wait for an SLA breach.

Q55. A maintenance window causes unexpected alerts to fire and page on-call staff. Prevention?

A: Schedule maintenance windows/downtime suppression in the monitoring tool ahead of time and confirm the affected devices are correctly tagged under that window.

Q56. How do you document a resolved incident for future reference?

A: Write a clear ticket/RCA covering the symptom, root cause, resolution steps, and preventive action, tagged so it's searchable in the knowledge base.

Q57. Bandwidth graphs show a steady climb toward saturation on a WAN link over several weeks. Action?

A: Proactively flag it for capacity planning/upgrade before it impacts users, and identify top talkers/applications via NetFlow/sFlow.

Q58. A P1 outage ticket sits untouched for 20 minutes because the assigned engineer is on break. Process fix?

A: Implement automated escalation timers so an unacknowledged P1 alert auto-escalates to the backup on-call after a set threshold.

Q59. Multiple monitoring tools report the same device outage at different severities. How do you reconcile?

A: Correlate through a single-pane-of-glass dashboard if available, or manually cross-reference; treat the highest reported severity as the ticket priority so urgency isn't downgraded.

Q60. You've fixed an issue but the monitoring dashboard still shows red. Next step?

A: Verify actual end-to-end service restoration (not just the specific check), then confirm or manually reset the monitoring tool's state — some require acknowledgment or a new polling cycle.

Q35. A site-to-site VPN drops whenever large file transfers start. Cause?

A: A possible MTU/fragmentation issue over the tunnel. Adjust MSS clamping or reduce MTU to prevent fragmented/dropped packets.

Q36. A user can VPN in from home but not from a hotel Wi-Fi. Cause?

A: The hotel network likely blocks VPN ports/protocols (UDP 500/4500) or requires a captive portal login first. Try a TCP-based VPN mode (SSL VPN over 443) instead.

Q37. After a VPN concentrator upgrade, users get "certificate not trusted" errors. Cause?

A: The new certificate wasn't issued by a trusted CA, or the client trust store hasn't been updated. Verify the certificate chain and push an updated trust store if needed.

Q38. Split-tunnel VPN users report both slow internet and slow internal resources. Cause?

A: Possible misconfigured split-tunnel routing sending all traffic through the tunnel (accidental full tunnel), saturating the concentrator. Verify the tunnel routing table.

Q39. Users behind VPN can reach servers by IP but not by name. Cause?

A: DNS isn't being pushed correctly by the VPN client, or split-DNS is excluding internal domains. Verify the DNS suffix and server push settings in the VPN profile.

Q40. A specific remote office consistently has poor VPN performance despite a good local speed test. Cause?

A: Possible ISP peering/congestion on the path to the VPN termination point, or asymmetric routing. Run MTR/traceroute from both ends to isolate the segment.

Wireless

Q41. Users on Wi-Fi experience random disconnects only in one area of the building. Cause?

A: A possible AP coverage gap, channel interference, or a rogue/interfering device. Run a site survey and check that AP's logs.

Q42. A new AP was added but clients still connect to a farther, weaker AP. Cause?

A: Client "sticky" roaming behavior, or AP power levels set too high/low causing improper cell sizing. Adjust transmit power and consider RSSI-based roaming assistance/band steering.

Q43. Wireless throughput is much lower than wired for the same user. What do you check?

A: Confirm the client's negotiated PHY rate/band (2.4 vs 5/6 GHz), channel utilization/interference, and AP uplink capacity. Also rule out too many clients on one AP.

Q44. Guest Wi-Fi users can access internal VLANs unexpectedly. Cause?

A: VLAN tagging misconfiguration on the AP or switchport, putting the guest SSID on the wrong VLAN. Verify the SSID-to-VLAN mapping and switchport trunk/access config.

Q45. Certain IoT devices fail to connect to a WPA2-Enterprise SSID. Cause?

A: The device likely doesn't support 802.1x/EAP properly. Consider a separate PSK or IoT-specific SSID/VLAN with MAC-based authentication.

Q46. After a wireless controller firmware update, multiple APs go offline. Cause?

A: A possible incompatible config push, CAPWAP control-plane failure, or a certificate mismatch between controller and AP. Check AP join status and logs.

Q47. Users near a microwave/kitchen area experience 2.4GHz Wi-Fi dropouts. Cause?

A: RF interference from the microwave overlapping 2.4GHz channels. Move affected users to 5GHz or shift the channel away from the interference.

A: Likely greylisting or a queue delay on the receiving mail server, or a DNS (MX/SPF) resolution issue causing retries. Check the mail queue and DNS record health.

Q97. FTP transfers fail in passive mode through the firewall but work in active mode (or vice versa). Cause?

A: The firewall isn't allowing the dynamic port range used by that FTP mode. Configure the firewall's FTP ALG or open the correct port range for the mode in use.

Q98. NTP sync failures are causing certificate and authentication errors across the network. Cause?

A: Clock drift beyond tolerance breaks time-sensitive TLS/Kerberos validation. Verify NTP server reachability and the correct time source configuration on all devices.

Q99. SNMP polling fails for a subset of devices after a security hardening project. Cause?

A: SNMP community strings/version changed, or an ACL now blocks the monitoring server's IP. Verify SNMP config matches on both the monitoring system and the devices.

Q100. An application server can be pinged and reached on some ports, but one particular TCP port randomly fails to connect. Cause?

A: Could be an intermediate firewall/ACL blocking that specific port, the application not listening or having crashed, or a stateful device dropping mid-session. Verify with telnet/nc and check server-side logs.